

# Containerized Bioinformatics Ecosystem for HPC

Yucheng Zhang<sup>1</sup>, Lev Gorenstein<sup>1</sup>, Payas Bhutra<sup>2</sup>, Ryan DeRue<sup>1</sup>

1. Rosen Center for Advanced Computing, Purdue University, West Lafayette, IN, USA

2. Department of Computer Science, Purdue University, West Lafayette, IN, USA

# Rationale

- Purdue RCAC have to manage multiple production systems, including 6 community clusters and ACCESS Anvil.
- Purdue has a large number of biological researchers studying various areas, such as agriculture, ecology, animal science, health science, etc.

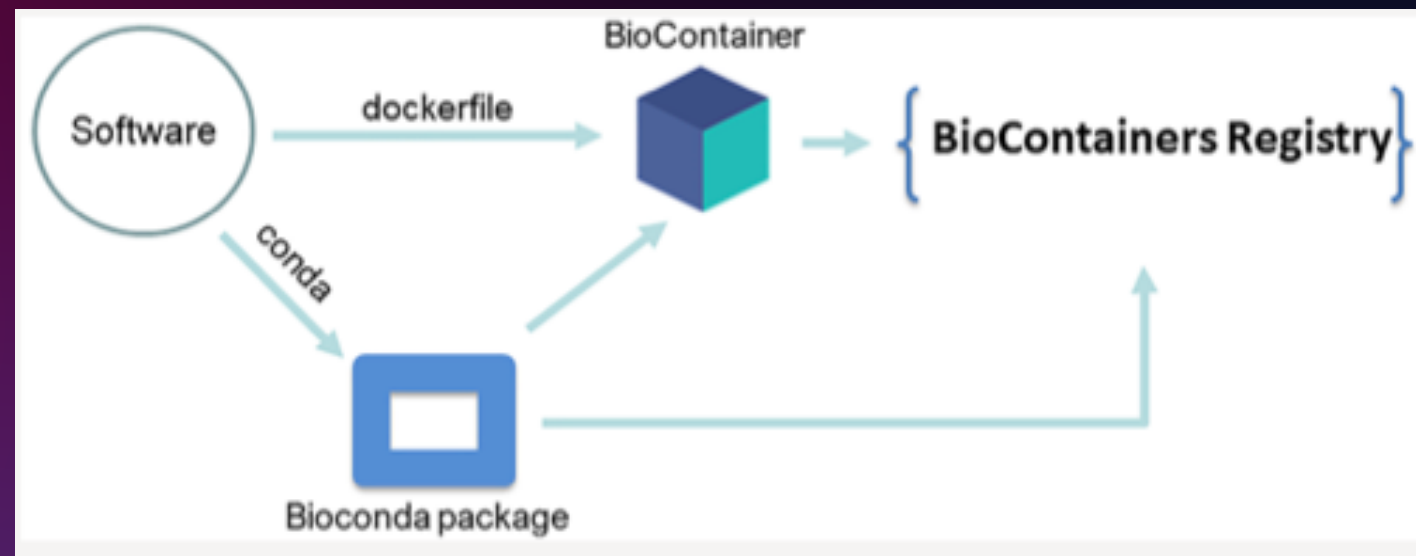
```
zhan4429@bell-fe00:~ $ module load trinity/2.12.0
Lmod has detected the following error: /depot/bioinfo/apps/modules/blast/2.2.26: (blast/2.2.26): child process exited abnormally
While processing the following module(s):
  Module fullname      Module Filename
  -----
blast/2.2.26           /depot/bioinfo/apps/modules/blast/2.2.26
seqclean/2011-02-22    /depot/bioinfo/apps/modules/seqclean/2011-02-22
PASA/r20140417         /depot/bioinfo/apps/modules/PASA/r20140417
trinity/2.12.0         /depot/bioinfo/apps/modules/trinity/2.12.0
```

**An easy and reliable approach to manage a large stack of bioinformatics applications is urgently needed.**

# BioContainers

## BIOCONDA®

- BioContainers is integrated with Bioconda, which is the conda channel for bioinformatics applications.
- BioContainers registry is the largest registry for bioinformatics applications.
- As of today, BioContainers provides containers for over 10 thousand bioinformatics applications.



J. Proteome Res. 2021, 20, 4, 2056–2061

# NGC container environment modules

NGC container environment modules are lightweight wrappers that make it possible to transparently use NGC containers as environment modules.

1. Allow HPC users to utilize familiar environment module commands.
2. Leverage all the benefits of containers, including portability and reproducibility.


<https://github.com/NVIDIA/ngc-container-environment-modules>

## Simplifying HPC Workflows with NVIDIA NGC Container Environment Modules

By Akhil Docca and Scott McMillan

Discuss (2) 0 Like

Tags: AI, Deep Learning, HPC / Supercomputing, machine learning, NGC, singularity





# Pull/build, test before deployment

Search applications from  
public registries

↓ Success

Pull container images to  
one HPC production system



Generate Lmod modulefiles




Submit sbatch jobs to test  
containers

↓ Success

Deploy modules to  
all HPC production systems

Failure →

Build our own container  
images



Failure

# Pulling images and generating Lmod modulefiles

## 1. `bioc_pull2sif.sh`

- a wrapper around “singularity pull”
- Outputting image names following the convention set by NGC container environment modules

## 2. `bioc_pull2mod.sh`

- Generate Lmod modulefile

## 3. `bioc_pull_and_module.sh`

- A wrapper combining the first two scripts
- Given a container URI, it will pull the container image and generate its modulefile

# Special Lmod modulefile setup

## 1. Add help/whatis information

## 2. GUI applications: bind X11 session information in ThinLinc

- `append_path("SINGULARITY_BIND", "/var/opt", ",")`
- `append_path("SINGULARITY_BIND", "/run/user", ",")`

## 3. Environment variables: environment variables associated with location to database or config files

- `pushenv("NAME", "value") ## set variable in host`
- `pushenv("SINGULARITYENV_NAME", "value") ## set variable inside container`

## 4. Adding executable path to PATH

- `pushenv("SINGULARITYENV_PREPEND_PATH", "/path/to/pkg/bin")`

## 5. Bind paths: bind database or config files

- `append_path("SINGULARITY_BIND", "hostdir:containerdir", ",")`



# Testing modules before deployment

```
zhan4429@bell-fe02:~ $ singularity exec abacas_1.3.1--pl5321hdfd78af_2.sif abacas.pl -r ref.fasta -q query.fasta -p nucmer
*****
* ABACAS: Algorithm Based Automatic Contiguation of Assembled Sequences *
* * * * *
* Copyright (C) 2008-10 The Wellcome Trust Sanger Institute, Cambridge, UK. *
* All Rights Reserved. *
* * * * *
*****

# Checking user options:
# -r Reference=ref.fasta
# -q Query=query.fasta
# -p nucmer
# -d 0 use sensitive mapping in nucmer i.e. --maxmatch
# Input checking done!!
PREPARING DATA FOR nucmer
delta-filter: error while loading shared libraries: libstdc++.so.6: cannot open shared object file: No such file or directory
show-tiling: error while loading shared libraries: libstdc++.so.6: cannot open shared object file: No such file or directory
Use of uninitialized value in addition (+) at /usr/local/bin/abacas.pl line 1001.
```

**Missing libraries**

Running analysis with real-world datasets is the only reliable way to make sure applications work as expected.



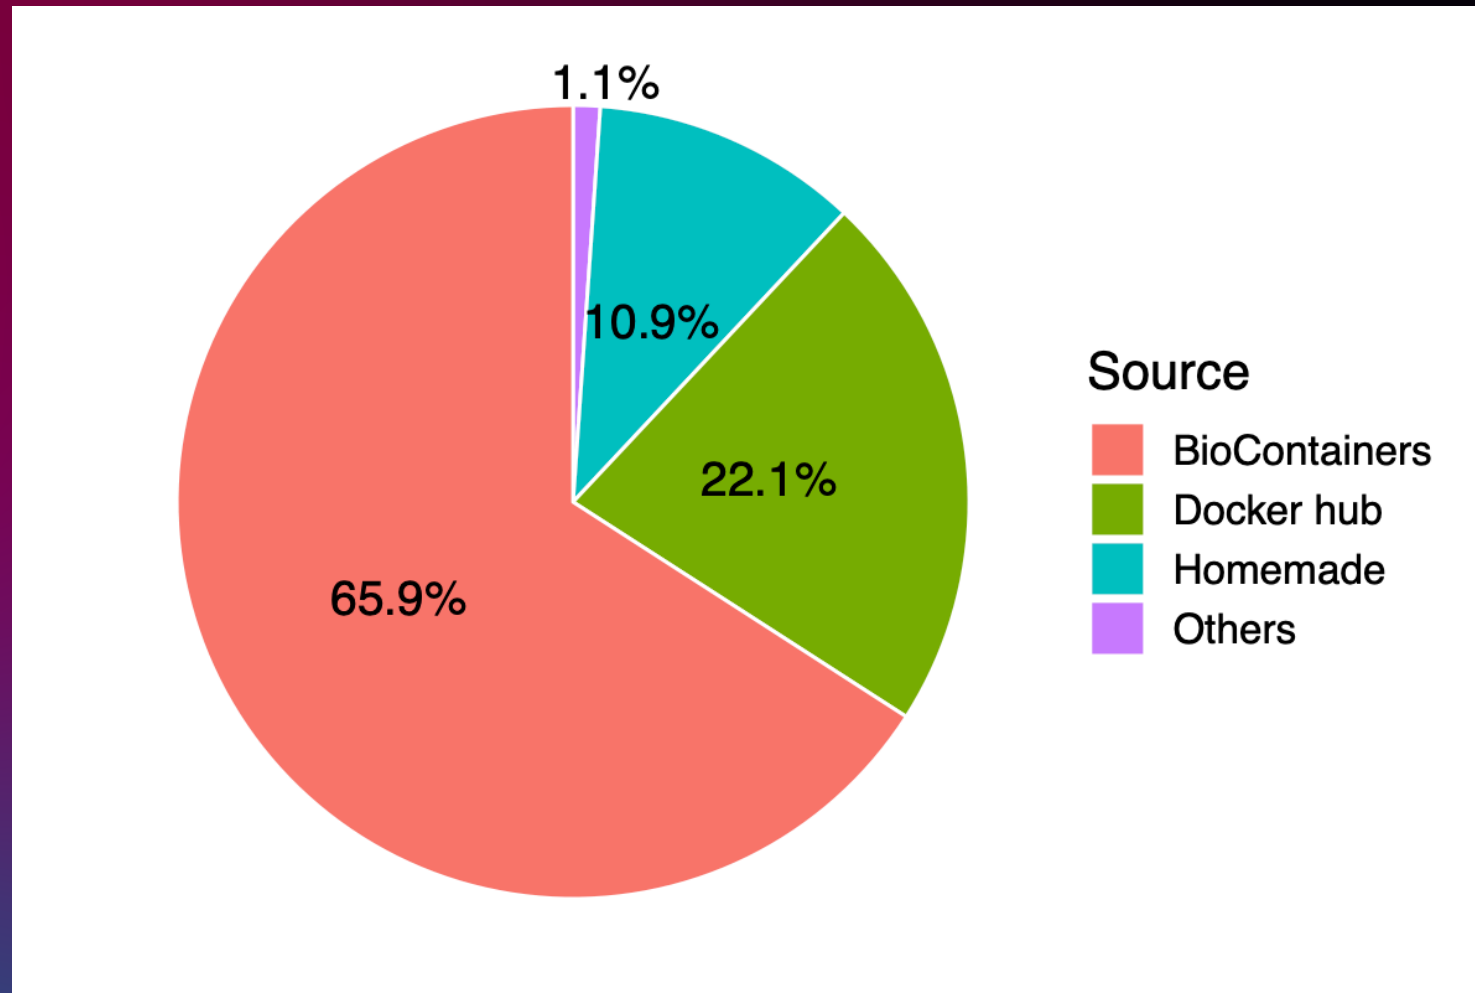


# 600 modules for 500 applications

BioContainers collection modules				
abacas/1.3.1		circexplorer2/2.3.8	hap.py/0.3.9	nextclade/1.10.3
abismal/3.0.0		circlator/1.5.5	helen/1.0	nextflow/21.10.0
abricate/1.0.1		circmpara2/0.1.2.1	hicexplorer/3.7.2	ngs-bits/2022_04
abyss/2.3.2		circos/0.69.8	hifiasm/0.16.0	ngsutils/0.5.9
abyss/2.3.4	(D)	ciriquant/1.1.2	hisat2/2.2.1	orthofinder/2.5.2
actc/0.2.0		clair3/0.1-r11	hmmer/3.3.2	orthofinder/2.5.4
advntr/1.4.0		clair3/0.1-r12	homer/4.11	(D)
afplot/0.2.1		clairvoyante/1.02	how_are_we_stranded_here/1.0.1	panacota/1.3.1
afterqc/0.9.7		clearcnv/0.306	htseq/0.13.5	panaroo/1.2.10
agat/0.8.1		clever-toolkit/2.4	htseq/1.99.2	pandaseq/2.11
alfred/0.2.5		clustalw/2.1	htseq/2.0.1	pandora/0.9.1
alfred/0.2.6	(D)	cnvkit/0.9.9-py	htslib/1.14	pangolin/3.1.20
alien-hunter/1.7.7		cnvnator/0.4.1	htslib/1.15	pangolin/4.0.6
alignstats/0.9.1		coinfinder/1.2.0	htslib/1.16	pangolin/4.1.2
allpaths-lg/52488		concoct/1.1.0	htstream/1.3.3	pangolin/4.1.3
alphafold/2.1.1		control-freec/11.6	humann/3.0.0	(D)
alphafold/2.2.0		cooler/0.8.11	hyphy/2.5.36	parallel-fastq-dump/0.6.7
alphafold/2.2.3	(D)	coverm/0.6.1	idba/1.1.3	parliament/2/0.1.11
amptk/1.5.4		crisprcasfinder/4.2.20	igv/2.11.9	parsnp/1.6.2
ananse/0.4.0		crispresso2/2.2.8	igv/2.12.3	pbmm2/1.7.0
anchorwave/1.0.1		crispresso2/2.2.9	impute2/2.3.2	pbtyper/1.0.4
angsd/0.935		crispresso2/2.2.10	infernal/1.1.4	pcangsd/1.10
angsd/0.937		crispritz/2.6.5	instrain/1.5.7	peakranger/1.18
angsd/0.939	(D)	cross_match/1.090518	instrain/1.6.3	pepper_deepvariant/r0.4.1
annogesic/1.1.0		crossmap/0.6.3	intarna/3.3.1	perl-bioperl/1.7.2-pl526
annovar/2022-01-13		csvtk/0.23.0	interproscan/5.54_87.0	phast/1.5
antismash/5.1.2		csvtk/0.25.0	iqtree/1.6.12	phd2fasta/0.990622
antismash/6.0.1		cufflinks/2.2.1	iqtree/2.1.2	phg/1.0
antismash/6.1.0	(D)	cutadapt/3.4	iqtree/2.2.0_beta	phrap/1.090518
anvio/7.0		cutadapt/3.7	isoseq3/3.4.0	phred/0.071220.c
anvio/7.1_main		cyvcf2/0.30.14	isoseq3/3.7.0	(D)
anvio/7.1_structure	(D)	dbg2olc/20180222	ivar/1.3.1	picard/2.25.1
any2fasta/0.4.2		dbg2olc/20200723	jcvi/1.2.7	picard/2.26.10
arcs/1.2.4		deepbgc/0.1.26	kaiju/1.8.2	picrust2/2.4.2
asgal/1.1.7		deepbgc/0.1.30	kallisto/0.46.2	picrust2/2.5.0
assembly-stats/1.0.1		deepconsensus/0.2.0	kallisto/0.48.0	(D)
atac-seq-pipeline/2.1.3		deepsignal2/0.1.2	khmer/3.0.0a3	pilon/1.24
ataqv/1.3.0		deeptools/3.5.1-py	kma/1.4.3	pindel/0.2.5b9
atram/2.4.3		deepvariant/1.0.0	kmc/3.2.1	pirate/1.0.4
atropos/1.1.17		deepvariant/1.1.0	kmer-jellyfish/2.3.0	pixy/1.2.7
atropos/1.1.31	(D)	delly/0.9.1	kneaddata/0.10.0	plasmidfinder/2.1.6
augur/14.0.0		delly/1.0.3	kover/2.0.6	platypus/0.8.1
augur/15.0.0	(D)	delly/1.1.3	kraken2/2.1.2	plink/1.90b6.21
augustus/3.4.0		delly/1.1.5	krakentools/1.2	plink2/2.00a2.3
augustus/3.5.0	(D)	diamond/2.0.13	kraken2/2.1.2	plotsr/0.5.4
bactopia/2.0.3		diamond/2.0.14	lambda/2.0.0	pomoxis/0.3.9
bali-phy/3.6.0		diamond/2.0.15	last/1268	popscle/0.1b
bam-readcount/1.0.0		dnaio/0.8.1	last/1356	pplacer/1.1.alpha19
bamgineer/1.1		dragonflye/1.0.13	last/1411	prinseq/0.20.4
			ldsc/1.0.1	prodigal/2.6.3
				prokka/1.14.6
				scvi-tools/0.16.2
				seidr/0.14.2
				sepp/4.5.1
				seqkit/2.0.0
				seqkit/2.1.0
				seqkit/2.1.0
				seqclean/1.10.09
				shasta/0.10.0
				shigeifinder/1.3.2
				shorah/1.99.2
				shortstack/3.8.5
				shovill/1.1.0
				sicer/1.1
				sicer2/1.0.3
				sicer2/1.2.0
				(D)
				signalp4/4.1
				signalp6/6.0-fast
				signalp6/6.0-slow
				(D)
				simug/1.0.0
				skewer/0.2.2
				slamdunk/0.4.3
				smoove/0.2.7
				snakemake/6.8.0
				snap-aligner/2.0.0
				snap/2013_11_29
				snaptools/1.4.8
				snippy/4.6.0
				snp-dists/0.8.2
				snp-sites/2.5.1
				snpeff/5.1d
				snpeff/5.1
				(D)
				snp genie/1.0
				snphylo/20180901
				snpsift/4.3.1t
				soapdenovo2/2.40
				sortmerna/2.1b
				sortmerna/4.3.4
				(D)
				souporcell/2.0
				sourmash/4.3.0
				sourmash/4.5.0
				(D)
				spaceranger/1.3.0
				spaceranger/1.3.1
				spaceranger/2.0.0
				(D)
				spades/3.15.3
				spades/3.15.4
				spades/3.15.5
				(D)
				sprod/1.0
				squeezemeta/1.5.1
				sra-tools/2.11.0-pl5262
				srst2/0.2.0



# Sources of container images

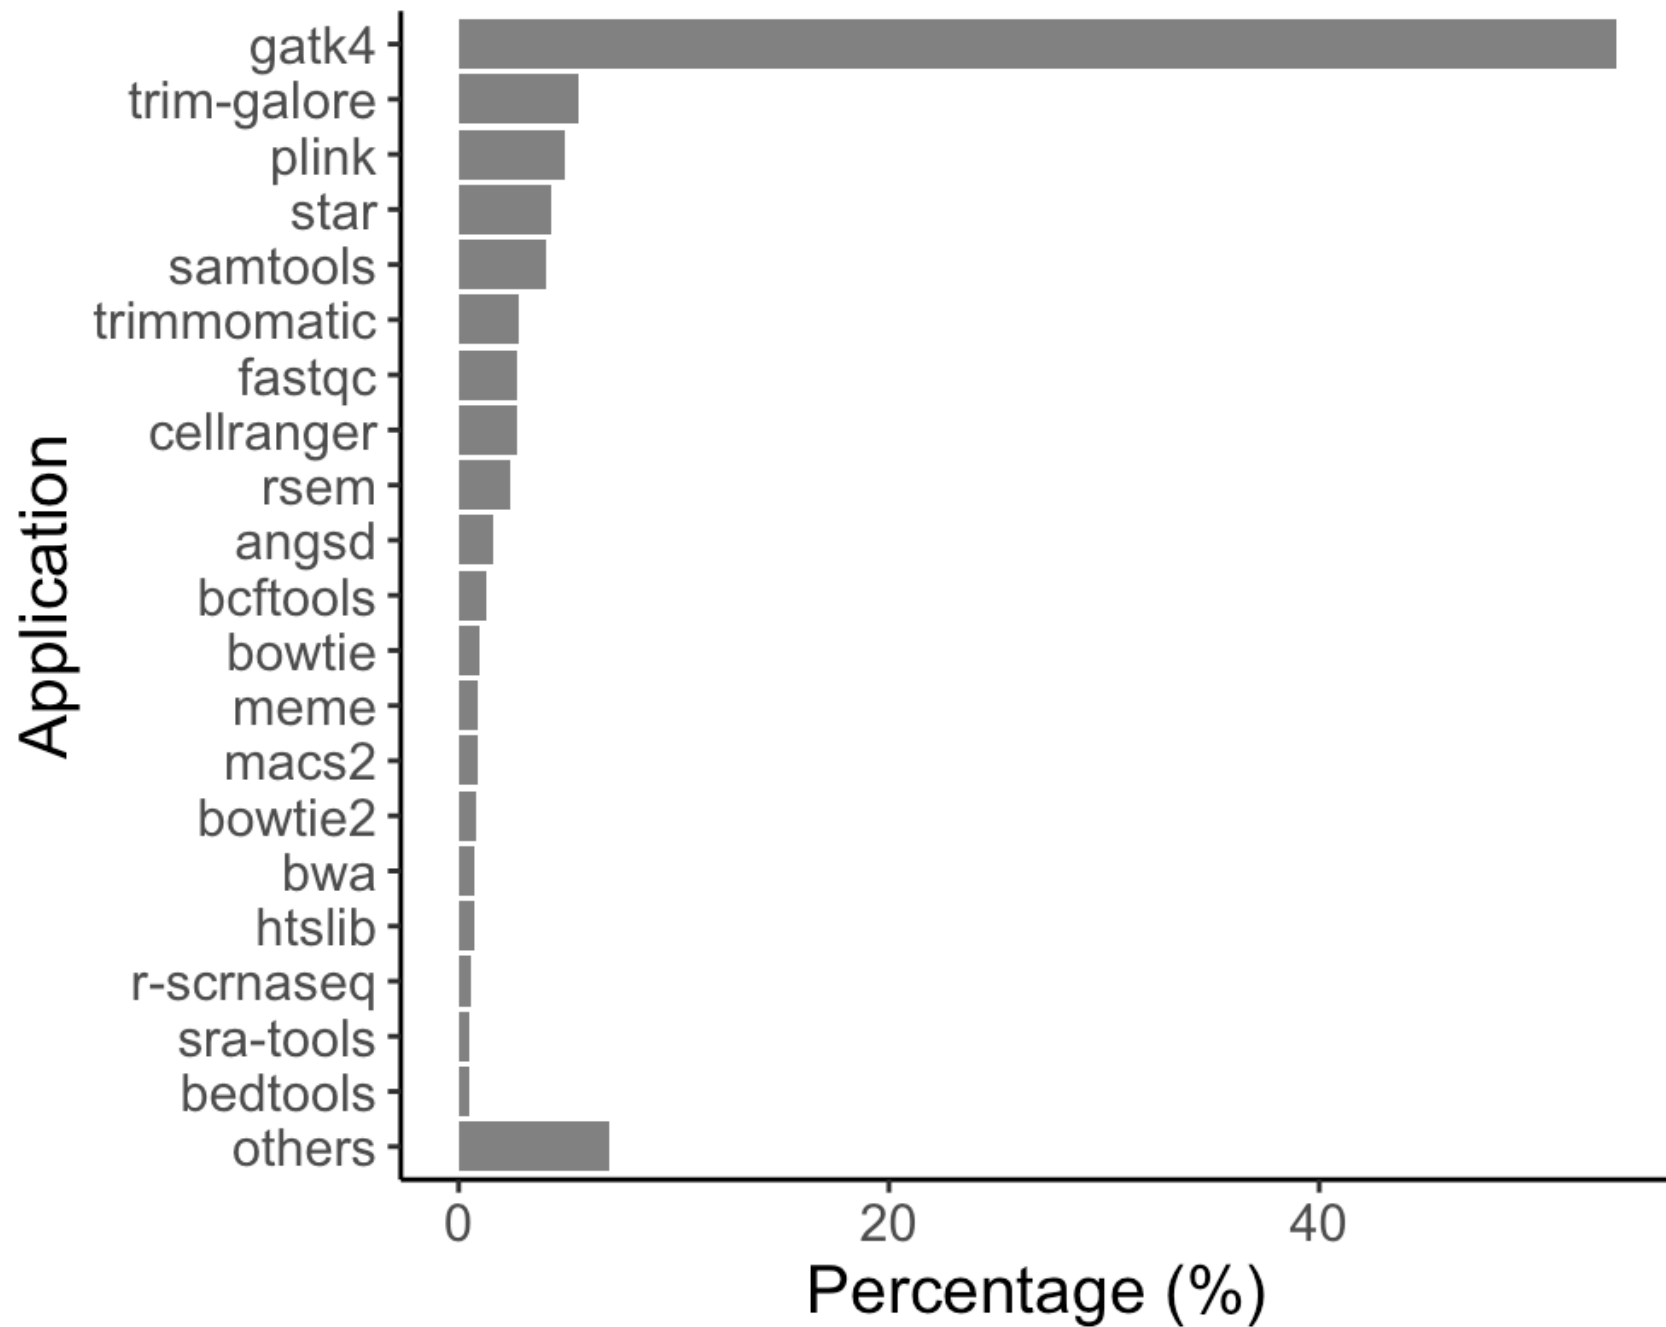


# Bioinformatics applications on HPC

1. **Fewer than 100:** most HPC centers
2. **~1000:** clusters designed for biosciences, e.g. NIH's Biowulf and Cornell's BioHPC
3. **Thousands:** the whole BioContainers project in TACC's clusters  
RollingGantryCrane (<https://github.com/TACC/rgc>)



**Too many is  
not necessary**





# Adding Jupyter support

```
Bootstrap: docker
```

```
From: quay.io/biocontainers/cellrank:1.5.1--pyhdfd78af_0
```

```
%labels
```

```
  Author: Yucheng Zhang zhan4429@purdue.edu
```

```
  Version: 1.5.1
```

```
%help
```

```
  CellRank with Jupyter support.
```

```
%post
```

```
  pip install ipython
```

```
  pip install ipykernel
```

Hopefully BioContainers developers can install `ipython` and `ipykernel` into all python-based container images.

# Open OnDemand Jupyter

```
{  
  "argv": [  
    "/usr/bin/singularity",  
    "exec",  
    "/apps/biocontainers/images/cellrank_1.5.1.sif",  
  ],  
  "python",  
  "-m",  
  "ipykernel_launcher",  
  "-f",  
  "{connection_file}"  
],  
  "display_name": "Cellrank",  
  "language": "python"  
}
```

`$HOME/.local/share/jupyter/kernels/`

The screenshot shows a Jupyter Notebook interface for Cellrank. The code in the notebook is as follows:

```
In [1]: import scvelo as scv  
import scanpy as sc  
import cellrank as cr  
import numpy as np  
  
scv.settings.verbosity = 3  
scv.settings.set_figure_params("scvelo")  
cr.settings.verbosity = 2  
  
In [2]: adata = cr.datasets.pancreas()  
scv.pl.proportions(adata)  
adata
```

The execution of the second cell shows a progress bar at 100% and a data visualization. The visualization consists of a pie chart and a stacked bar chart. The pie chart shows 81% spliced and 19% unspliced. The stacked bar chart shows the proportions of spliced and unspliced cells across different clusters.

clusters	spliced	unspliced
Ngn3 low EP	88%	12%
Ngn3 high EP	85%	15%
Fev+	79%	21%
Beta	79%	21%
Alpha	80%	20%
Delta	81%	19%
Epsilon	82%	18%

# Bundle applications into a single container image

With containers, it is easy to install not just a single application, but also bundles and collections of multiple applications working in concert and dedicated to a specific research workflow.



**R-RNAseq**  
Customized R container for RNAseq analysis.

- ComplexHeatmap
- DESeq2
- DEXSeq
- edgeR
- ggrepel
- Limma
- pheatmap
- tidyverse



**R-scRNAseq**  
Customized R container for scRNAseq analysis.

- CoGAPS
- DESeq2
- doSNOW
- DropletUtils
- edgeR
- Limma
- miQC
- monocle
- monocle3
- Nebulosa
- rliker
- scCATCH
- scDbfFinder
- SCHNAPPs
- scMappR
- seurat
- seurat-wrappers
- SingleR
- SnapATAC
- SoupX
- tidyverse
- tricycle
- velocity.R
- And more



# Open OnDemand

- For bioinformatics applications that use a native graphics user interface (GUI) and that have a large computational or memory footprint, we employ Open OnDemand to allow users to easily allocate appropriate amount of resources and submit jobs through a convenient web interface.
- We create a simple workflow for rapid deployment of containers to Open OnDemand of any cluster.
  1. **default\_biocontainer\_template**: a template Open OnDemand application directory for a generic VNC desktop application.
  2. **deploy\_biocontainer**: a helper script that makes a copy of the template directory and performs the necessary substitutions to the relevant files.

The screenshot displays the Open OnDemand web interface for configuring the MaxQuant application. The breadcrumb navigation at the top reads "Home / My Interactive Sessions / MaxQuant".

**Bioinformatics Apps**

- Interactive Apps
  - CryoSPARC
  - Integrative Genomics Viewer
  - MEGAN
  - MaxQuant** (selected)
  - QualiMap
  - Rstudio\_scrNaseq
  - Tassel5

**Interactive Apps**

- Desktops
  - Desktop
- GUIs
  - MATLAB

**MaxQuant**

This app will launch MaxQuant on the [Anvil](#) cluster.

Allocation: asc170016 (81643.7 SUs remaining)

Queue (partition): shared

- GPU-only allocations MUST use the 'gpu' queue
- CPU-only allocations MAY NOT use the 'gpu' queue

Wall Time in Hours: 1

Number of nodes: 1

Node count is limited to 1 on 'shared' queue

Cores: 2

Number of cores (up to 128) for a shared job. Non-shared jobs will have exclusive nodes and be charged at 128 cores per node requested



# Turning containers into OOD apps

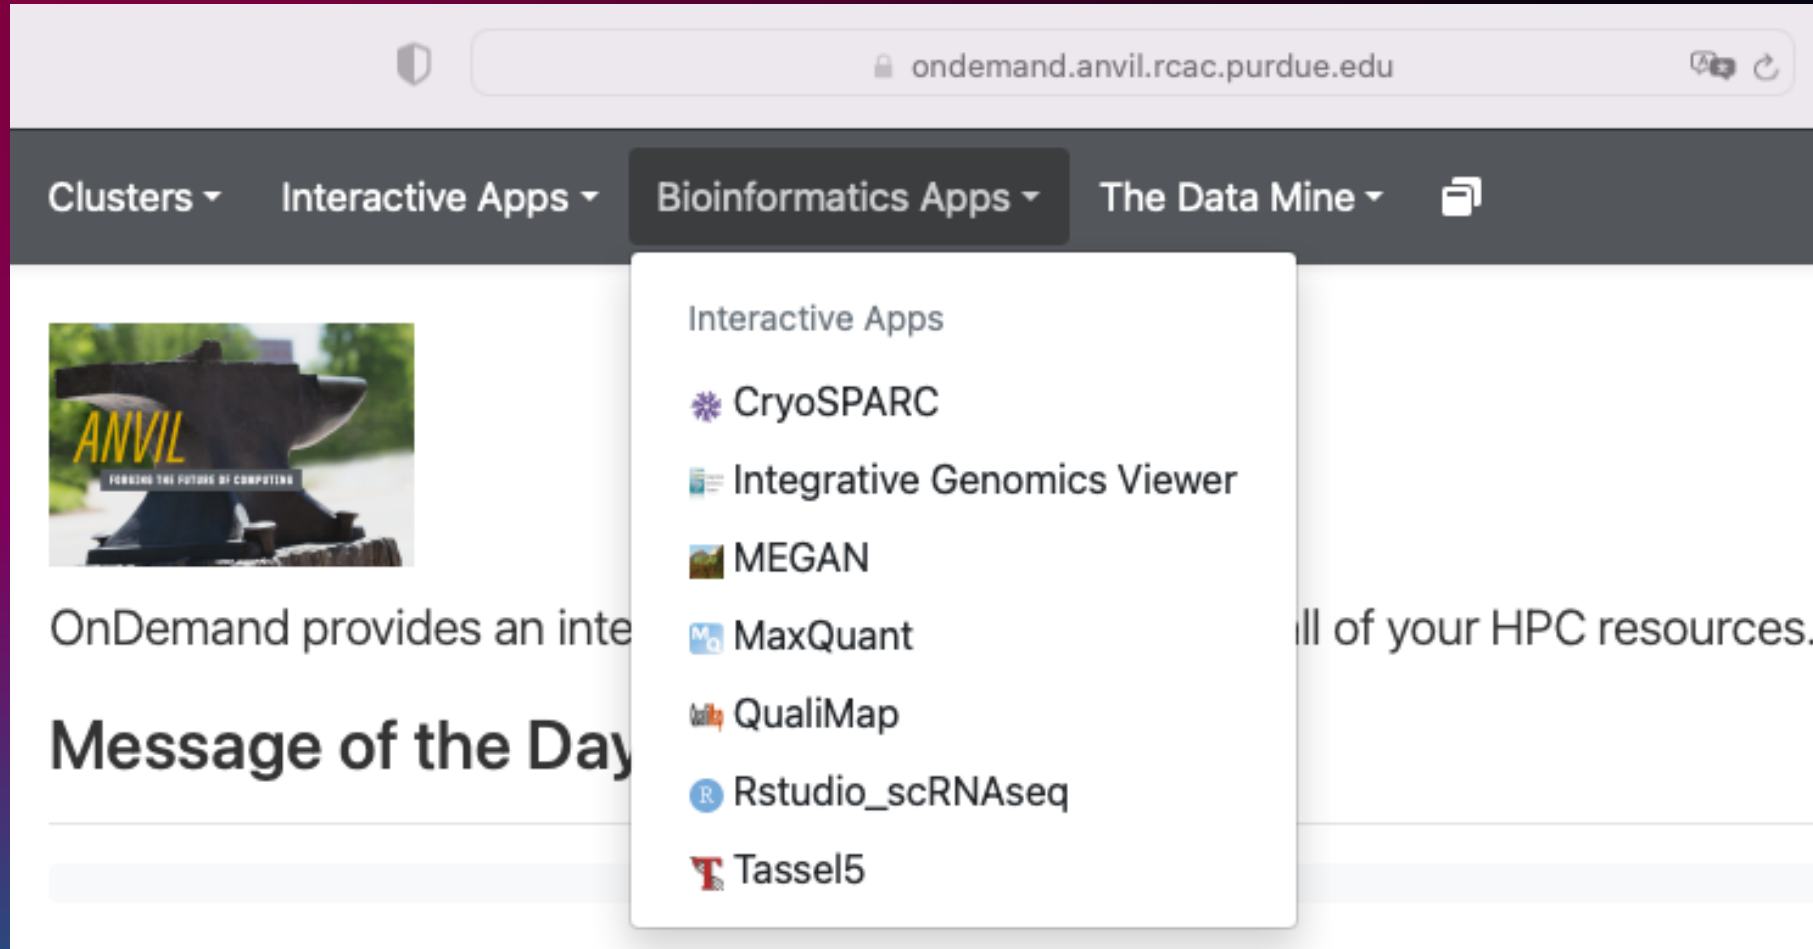
With the template directory `default_biocontainer_template` and the `deploy_biocontainer` script, we can easily turn container images into Open OnDemand interactive applications using a one-line command:

```
deploy_biocontainer --name appName \  
                    --directory folderName \  
                    --image app.sif \  
                    --command launchCommand \  
                    default_biocontainer_template
```

```
ENTRY="[Desktop Entry]  
Type=Application Name=IGV  
Comment=  
Exec=/bin/bash -lc \"singularity exec /apps/biocontainers/images/igv_2.12.3.sif igv.sh\"  
Path=  
Terminal=false  
StartupNotify=false  
Categories=Cluster\"  
echo -e \"$ENTRY\" > \"${AUTOSTART}/igv.desktop\"
```

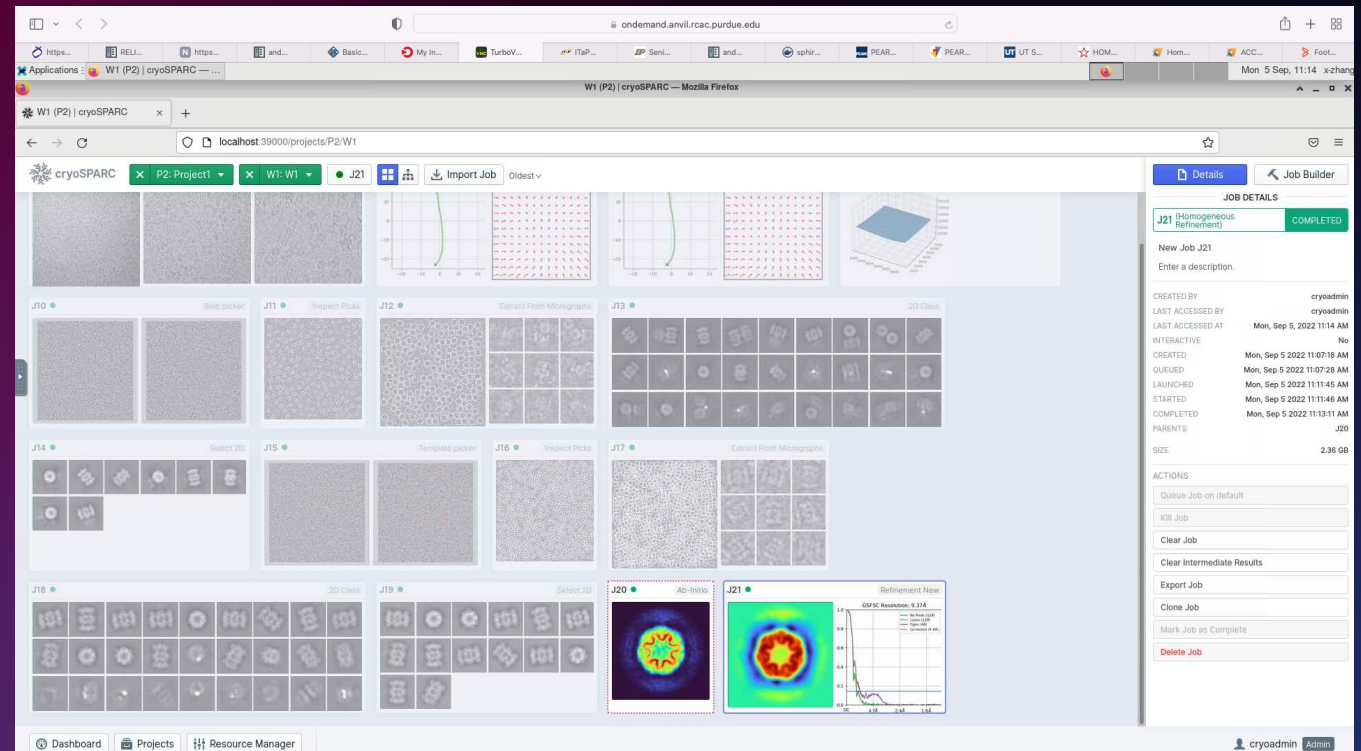
A snippet from `xfce.sh` for the genomic browser IGV

# Open OnDemand Applications



The screenshot shows a web browser window with the URL `ondemand.anvil.rcac.purdue.edu`. The navigation bar includes "Clusters", "Interactive Apps", "Bioinformatics Apps", and "The Data Mine". The "Bioinformatics Apps" menu is open, displaying a list of applications: Interactive Apps, CryoSPARC, Integrative Genomics Viewer, MEGAN, MaxQuant, QualiMap, Rstudio\_scRNAseq, and Tassel5. The background of the page features an image of an anvil with the text "ANVIL FORGING THE FUTURE BY COMPUTING" and a "Message of the Day" section.

# Not only bioinformatics



Thank **Michael Dickens** from Texas A&M University for showing us how to build and run the cryoSPARC container.

# Helper command

**Note**

Since `BRAKER` is a pipeline that trains `AUGUSTUS`, i.e. writes species specific parameter files, BRAKER needs writing access to the configuration directory of AUGUSTUS that contains such files. This installation comes with a stub of AUGUSTUS configuration files, but you **must** copy them out from the container into a location where you have write permissions.

A helper command `copy_augustus_config` is provided to simplify the task. Follow the procedure below to put the config files in your scratch space:

```
$ mkdir -p $RCAC_SCRATCH/augustus
$ copy_augustus_config $RCAC_SCRATCH/augustus
$ export AUGUSTUS_CONFIG_PATH=$RCAC_SCRATCH/augustus/config
```

**Warning**

Using `#!/bin/sh -l` as shebang in the slurm job script will cause the failure of some biocontainer modules. Please use `#!/bin/bash` instead.

To run SortMeRNA on our clusters:

```
#!/bin/bash
#SBATCH -A myallocation # Allocation name
#SBATCH -t 1:00:00
#SBATCH -N 1
#SBATCH -n 1
#SBATCH --job-name=sortmerna
#SBATCH --mail-type=FAIL,BEGIN,END
#SBATCH --error=%x-%j-%u.err
#SBATCH --output=%x-%j-%u.out

module --force purge
ml biocontainers sortmerna

sortmerna --ref silva-bac-16s-id90.fasta,silva-bac-16s-db \
--reads set2_environmental_study_550_amplicon.fasta \
--fastx --aligned Test
```

RCAC Biocontainers

latest

Search docs

FREQUENTLY ASKED QUESTIONS

Frequently Asked Questions

SINGULARITY

Singularity

APPLICATION LIST

- Abacas
- Abismal
- Abricate
- Abys
- Actc
- Advntr
- Afplot
- Afterqc
- Agat
- Alfred
- Alien-hunter
- Alignstats
- Allpathsig
- Alphafold
- Amptk
- Ananse
- Anchorwave
- ANGSD
- Annogesic
- ANNOVAR
- Antismash
- Anvio

biocontainer-doc.readthedocs.io

RCAC Biocontainers documentation!

Edit on GitHub

## RCAC Biocontainers documentation!

This is the user guide for biocontainer modules deployed in Purdue High Performance Computing clusters. More information about our center is available here (<https://www.rcac.purdue.edu>).

If you have any question, contact me(Yucheng Zhang) at: [zhan4429@purdue.edu](mailto:zhan4429@purdue.edu)

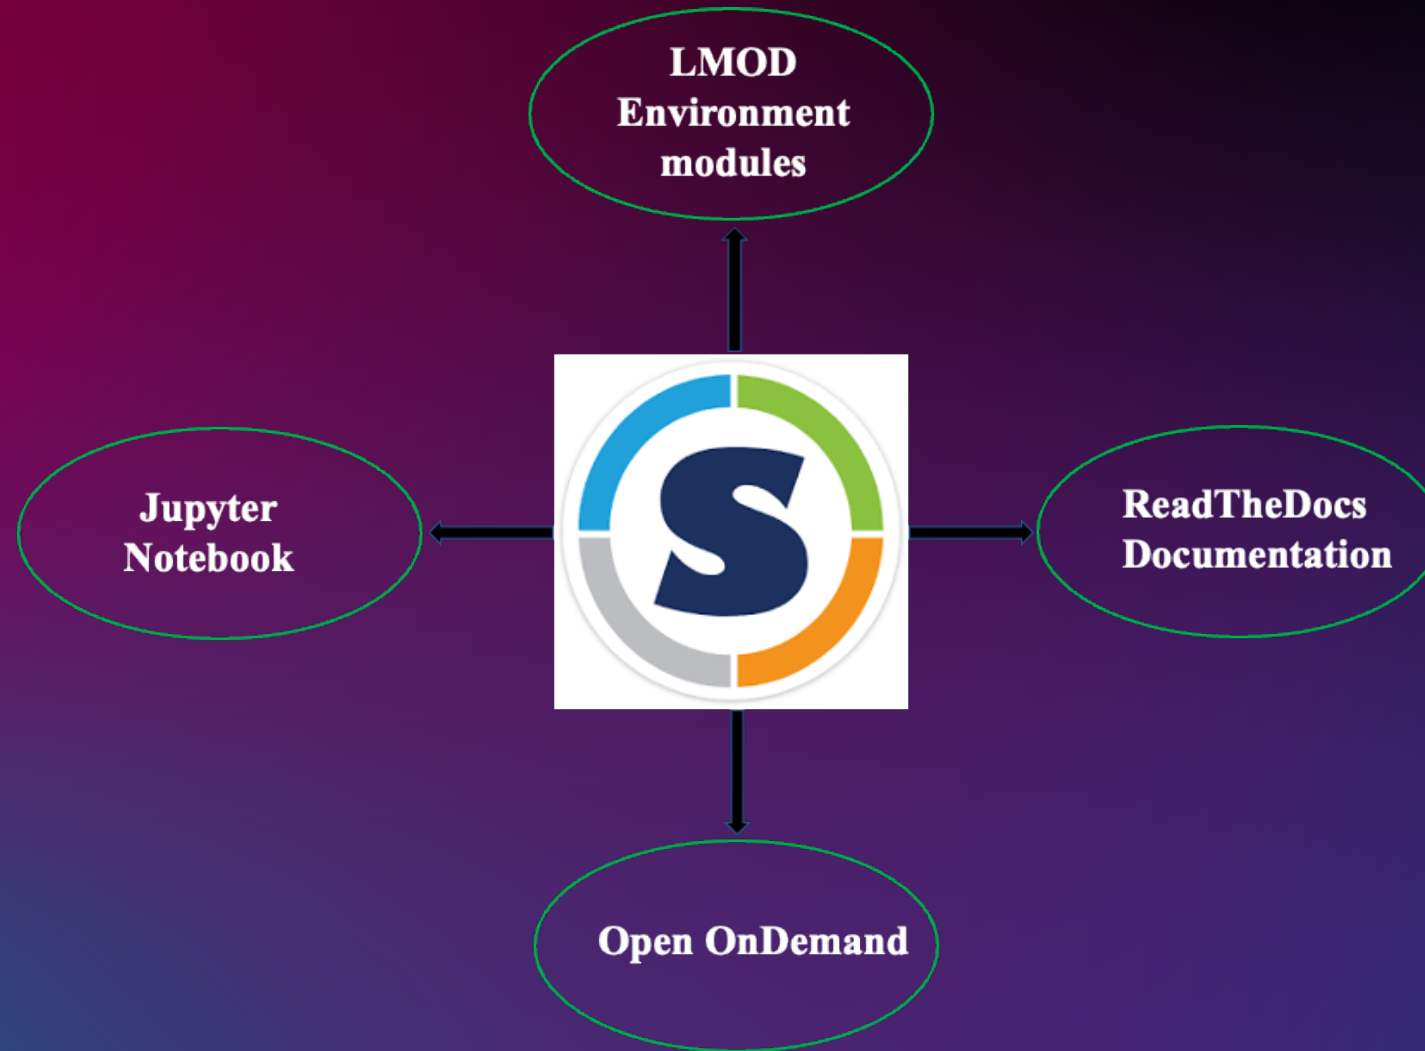
**Warning**

Do not use both `bioinfo` and `biocontainers` in your job script, because loading `bioinfo` will cause the failure of loading many modules including `biocontainers` in `Brown`, `Halstead`, `Scholar`, `Workbench`, and `Gilbreth`. Since RCAC will not provide support to `bioinfo` in the future clusters, we recommend users to just use `biocontainers`.

<https://biocontainer-doc.readthedocs.io/en/latest/>



# Containerized Bioinformatics Ecosystem



# Interested in building a similar ecosystem in your center?

Search or jump to... Pull requests Issues Marketplace Explore

PurdueRCAC / Biocontainers Public

Code Issues Pull requests Actions Projects Wiki Security Insights Settings

main 1 branch 0 tags

Go to file Add file Code

Yucheng Zhang add pull and module scripts dd406a6 23 seconds ago 35 commits

definition_files	add module files	2 days ago
docs	mv cat.py	14 hours ago
module_files	add module files	2 days ago
scripts	add pull and module scripts	23 seconds ago
README.md	Update README.md	last month

README.md

## Purdue RCAC Biocontainer

Since 2021, Purdue Research Computing team began to deploy containerized scientific applications on Purdue HPC clusters as well as XSEDE Anvil. Among them, **Biocontainer** is a collection of applications for bioinformatic analysis. Thanks to the portability of containerized applications, biocontainer modules can be easily deployed on Anvil and Purdue RCAC's 6 clusters. The number of biocontainer applications grows fast. Within 6 months since we initiated this project, the number of fully tested and deployed bioinformatic applications exceeded 300, and is estimated to reach 600-800 by the end of 2022.

It is worth to mention that before deploying the modules into production, system administrators also use them to run real-world datasets to make sure the applications work as expected. This is a time-consuming step, but it is quite essential, because our tests revealed that some public containers failed to work due to various kinds of issues. For such flawed containers, we built our own container images from scratch instead.

About  
Containerized bioinformatics applications deployed in Purdue community clusters and Access Anvil

Readme  
2 stars  
1 watching  
0 forks

Releases  
No releases published  
[Create a new release](#)

Packages  
No packages published  
[Publish your first package](#)

Languages

Lua	95.2%	Perl	3.3%
Shell	1.1%	Roff	0.4%

## Contributions are welcome!

- If you find issues or bugs, please open an issue in the GitHub repository.
- Our goal is to improve the Biocontainer project together with all centers, and we need your ideas and input to keep on improving.
- We welcome any contribution including scripts, modulefiles, definition files, etc.

git clone <https://github.com/PurdueRCAC/Biocontainers.git>

# Thank you!

## Contributors

- Yucheng Zhang
- Lev Gorenstein
- Payas Bhutra
- Ryan DeRue

## Purdue RCAC

- Preston Smith
- Xiao Zhu (Intel)
- Arman Pazouki
- Amiya Maji
- Geoffrey Lentner
- Guangzhen Jin
- Sarah Rodenbeck
- Steve Kelley
- Yang Hong
- Tsai-wei Wu
- Nannan Shan
- Eric Adams
- Ruyi Li

# HUST22 Committee

